A COMPARISON OF CLASSIFIER PERFORMANCE FOR VIBRATION-BASED TERRAIN CLASSIFICATION

Eric Coyle* and Emmanuel G. Collins, Jr.
Department of Mechanical Engineering
FSU-FAMU College of Engineering
Tallahassee FL, 32310
(coyleer, ecollins@eng.fsu.edu)

ABSTRACT

The ability to recognize the encountered terrain is an essential part of any terrain-dependent control system designed for mobile robots. Terrains such as sand and gravel make vehicle mobility more difficult and thus reduce vehicle performance. To alleviate this problem the vehicle control system can be tuned for maximum speeds, turning angles, accelerations and other conditions to help adapt to various terrains. Terrain classification can be used to automate the switch from one control mode to another. This paper compares the performance of several classifiers on the problem of vibration-based terrain classification. The purpose of this comparison is to assess the strengths and weaknesses of these techniques in order to better understand the tools available in developing future vibration-based terrain classification algorithms.

1. INTRODUCTION

Development of intelligent systems that improve autonomous ground vehicle (AGV) performance on difficult off-road terrains is important for military missions. Terrain Response, originally designed for the Land Rover LR3 SUV is an example of a terrain-dependent control system (Vanderwerp, 2005). This system, which has since been implemented on other Land Rover vehicles such as the Freelander, has several driving modes that adjust vehicle parameters such as anti-lock braking, throttle response, and differential locking based on predefined settings. The determination of when to switch between these modes is left up to the driver of the vehicle. However, without automated terrain recognition such a system could never be implemented on an autonomous military vehicle. Thus, terrain dependent control systems for autonomous vehicles require the ability to recognize the underlying surface as well as implement appropriate vehicle control system adjustments. There are two general types of sensors that have proven effective for terrain detection, vision sensors and vibration sensors.

Vision-based terrain detection requires the use of cameras or laser range finding sensors, also called LADAR. Using LADAR systems a 3D map of the environment can be obtained, which can then be used to classify the shapes into groups of vegetation, shrubbery and trees (Vandapel et al., 2006). Other research uses terrain maps created using LADAR to determine whether encountered surfaces are navigable or non-navigable (Wolf et al., 2005). Additionally, similar terrain maps can be computed using stereo imagery as in (Se et al., 2005). Other vision-based research has sought to use cameras to characterize the roughness, slope, discontinuity and hardness of the terrain in hopes of using these characteristics to navigate the terrain more appropriately (Howard and Seraji, 2001). More recent research has shown that many terrains can be classified based on the observed color in camera images (Manduchi et al., 2005).

Terrain classification using vibration sensors has been demonstrated using a variety of techniques. Recent research in vibration-based terrain classification, which was originally suggested in (Iagnemma and Dubowsky, 2002), has proven that vibration signals possess terrain signatures when transformed into the frequency domain using a Fast Fourier Transform (FFT) (Sadhukhan and Moore, 2003), (DuPont et al., 2005a), (Ojeda et al., 2006), (Weiss et al., 2006) , (DuPont et al., 2008b). The research in (DuPont et. al., 2005a) has also shown improved accuracy by incorporating multiple vibration measurements recorded using an Inertial Measurement Unit (IMU). Additionally, research has shown that the use of eigenspace feature extraction and selection through Principal Component Analysis (PCA) can improve both accuracy and classification time (DuPont et al., 2005b), (Brooks et al., 2005), (DuPont, et al., 2006). A detailed explanation of current terrain classification techniques as well as a basis for why such techniques work can be found in (Dupont, et al., 2008a). Ideally a robust terrain recognition algorithm will rely on both vision and vibration-based terrain classification systems, just as a human driver uses vision and feel to determine the terrain type. Although some research on fusing the two types of classification has begun on planetary rovers (Halatci et al., 2007), at this time each method could benefit from additional testing and algorithm refinement.

The general approach of using frequency based features and PCA for vibration-based classification is fairly

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
DEC 2008	**N/A**	**-**

4. TITLE AND SUBTITLE	
A Comparison Of Classifier Performance For Vibration-Based Terrain Classification	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Department of Mechanical Engineering FSU-FAMU College of Engineering Tallahassee FL, 32310**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM002187. Proceedings of the Army Science Conference (26th) Held in Orlando, Florida on 1-4 December 2008, The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**UU**	**7**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

standard, but published research has shown significant variation in the choice of classifier. While (Weiss et al., 2006) used a support vector machine classifier with a radial basis kernel function, in (Brooks et al., 2005) a discriminant function classifier is used in a one against one scheme. In (Sadhukhan and Moore, 2003), (DuPont et al., 2005b), (DuPont, et al., 2006), and (DuPont, et al., 2008b) a probabilistic neural network (PNN) classifier is used for vibration-based terrain classification and neural networks were used in (Ojeda et al., 2006) to test the different sensor modalities. With so many classification techniques, a comparison of published techniques was given in (Weiss et al., 2007) to try and determine the most appropriate method of vibration-based terrain classification. This work included classifiers based on decision trees, naïve Bayes and K-nearest neighbor to go along with the works of (Brooks et al., 2005), (Weiss et al., 2006) and (Sadhukhan and Moore, 2003), the latter of which preceded the work of (DuPont et al., 2005b). This comparison ultimately concluded that a support vector machine classifier using a radial basis function as a kernel function was the most appropriate technique. However, this research did not make use of PCA for feature selection and extraction, a process that can be vital to high performance with specific classifiers. Nor did the comparison in (Weiss et al., 2007) make use of pitch rates and role rates, instead relying solely on the vehicle acceleration perpendicular to the ground surface, also referred to as the vertical acceleration. Additionally, this research was focused on finding the best approach to terrain classification among the previously published techniques, rather than studying the strengths and weaknesses of traditional pattern recognition techniques as it relates to vibration-based terrain classification.

This paper concentrates on analyzing the benefits and drawbacks of applying several different classifiers, typically used in statistical pattern recognition, when applied to vibration-based terrain classification. It is not expected that a single classifier will necessarily stand out as the best classifier, rather each classifier is expected to display both strengths and weaknesses as it relates to vibration-based terrain classification. A pattern recognition technique called cross validation (Duda et al., 2001) is used for tuning both the classifiers and PCA. Cross validation is intended to determine appropriate classifier parameters while avoiding the problem of over-training, which is characterized by good performance on the test set, but poor performance on the overall population. Many published works in terrain classification have failed to completely address this problem of over training. Metrics of classification time, *mean accuracy,* which is the average of the accuracies on each terrain and *minimum accuracy,* which is the lowest accuracy of any terrain, are used to convey the advantages and deficiencies of each classifier. The combination of these two accuracies should indicate whether a classifier not only yields high classification accuracy but also indicate if a classifier can handle a larger variety of terrains. Additionally, difficulty in training will be addressed, but to a lesser extent than the accuracies and classification time.

2. CLASSIFIERS FROM STATISTICAL PATTERN RECOGNITION

In pattern recognition, a classifier is generally trained based on an observed set of p patterns $\boldsymbol{T} = \{\boldsymbol{t_1} \quad \boldsymbol{t_2} \quad \dots \quad \boldsymbol{t_p}\}$ referred to as the training set. Here, each training pattern $\boldsymbol{t_m}$ has k features, i.e. $\boldsymbol{t_m} \in \Re^k$ and corresponds to one of c classes in the set of classes $\boldsymbol{\omega} = \{\omega_1 \quad \omega_2 \quad \dots \quad \omega_c\}$. The classifier then attempts to classify a new pattern of unknown class, called a test pattern $\boldsymbol{x}$ with k features, as belonging to the best choice of c classes. The uniqueness of statistical classifiers is in how $\boldsymbol{T}$ is used to determine the choice of class in $\boldsymbol{\omega}$ for $\boldsymbol{x}$.

Classifiers in statistical pattern recognition generally fall into five categories. These categories are probabilistic methods, discriminant function analysis, nearest neighbors, decision trees and neural networks. The following subsections will give a brief description of each of these categories and the individual classifiers from these categories chosen to be included in the vibration-based terrain classification comparison. For a more detailed description of these and other pattern recognition methods see (Duda et al., 2001).

2.1 Probabilistic Methods

Probabilistic classifiers are largely based on Bayes decision rule. This rule states:

$$\text{if } p(\omega_i|\boldsymbol{x}) > p(\omega_j|\boldsymbol{x}) \text{ for all } j \neq i, \tag{1}$$

then $\boldsymbol{x}$ most likely belongs to class ω_i. Although they infrequently occur, ties can be broken arbitrarily since any class with the same probability is considered equally likely to be the correct class. Probabilistic based classifiers estimate probability distributions for each class using either a parametric or nonparametric approach. A parametric estimation technique assumes a given form of the probability distributions, while a nonparametric technique does not require such an assumption.

Both a parametric and a nonparametric technique will be considered for the vibration-based terrain classification comparison. Maximum likelihood estimation is one of the two most commonly used parametric estimation techniques in pattern recognition. It estimates the distribution parameters by choosing the parameters that make the training data "most likely" to be observed. This paper will consider maximum likelihood estimation using

an assumed Gaussian distribution. Since it is a non-parametric technique, Parzen window estimation can be used to accurately estimate any smooth distribution. The window function, which defines the influence of each training sample, used in this paper is the commonly used Gaussian function. It should be noted that this implementation of Parzen window estimation is extremely similar to a (PNN) classifier like the ones used in (Sadhukhan and Moore, 2003), (DuPont et al., 2005a), (DuPont et al., 2005b), (DuPont et al., 2006) and (DuPont et al., 2008b). The merits of Parzen window estimation will be compared to maximum likelihood estimation to represent the benefits of parametric and nonparametric probability estimation.

2.2 Discriminant Functions

Discriminant functions are used to determine appropriate relationships between the k feature variables and appropriate decision boundaries. Generally speaking, the idea behind discriminant functions is to determine a function that will appropriately represent the shape of the decision boundaries and then estimate the unknown parameters, i.e. coefficients, in the discriminant function. For instance, consider a set of linear discriminant functions that assume boundaries based on separating hyperplanes

$$g_i(\boldsymbol{x}) = \boldsymbol{a}_i \boldsymbol{x} + a_{io} \tag{2}$$

where $g_i(\boldsymbol{x})$ is the discriminant function for the i^{th} class, $\boldsymbol{a}_i$ is a weight vector and a_{io} is a bias term. Here, the weights $\boldsymbol{a}_i$ and bias a_{io} are determined from $\boldsymbol{T}$ using computer learning and a proposed decision rule such as

$$\text{assign } \boldsymbol{x} \text{ to } \omega_i \text{ if } g_i(\boldsymbol{x}) > g_j(\boldsymbol{x}) \text{ for all } j \neq i \tag{3}$$

or

$$\text{assign } \boldsymbol{x} \text{ to } \omega_i \text{ if } g_i(\boldsymbol{x}) > 0\,. \tag{4}$$

Probabilistic methods can in fact be viewed as a form of discriminant functions, where the relationship between feature variables and the decision boundaries are probabilistic in nature. One problem with discriminant functions is that many learning methods for discriminant functions are designed to solve a two class problem. Since terrain classification is unlikely to be a two class problem, it will likely require the use of a one against one or a one against the rest decision scheme. However, both schemes can create what is known as ambiguous regions where several classes are deemed to be possible as the true class.

There are methods that bypass the problem of ambiguous regions, including Kesler's construction, which is a form of linear discriminant analysis. Physically speaking, Kesler's construction partitions the feature space into regions corresponding to the individual classes, where the boundaries between regions are made up of hyperplanes. Problems with Kesler's construction tend to occur most often when separating hyperplanes are poor choices for the shape of the decision boundaries. Support vector machines, which are another form of linear discriminant analysis, can sometimes solve this problem by first mapping the original features to a higher dimensional space. The idea is that in this new space, separating hyperplanes may be more appropriate. Additionally, support vector machines maximize the margin between the separating hyperplane and the closest training points, which are called support vectors, resulting in what is known as the optimal hyperplane. Since this optimization problem may not always be solvable, that is a separating hyperplane may not exist in the higher dimensional space, the choice of kernel function is highly important. Both Kesler's construction and vote based one against one support vector machines are used in the comparison presented later in this paper. The SVM algorithm used in this paper is the LIBSVM algorithm (Chang and Lin, 2005).

2.3 Nearest Neighbors

Perhaps the simplest and most intuitive category of classifiers is that of nearest neighbor based classifiers. A nearest neighbor based classifier attempts to classify a test pattern $\boldsymbol{x}$ based on the class of the "closest" training pattern or patterns. However, the obvious question this invites is how to determine which training patterns in $\boldsymbol{T}$ are the closest to the test pattern. The Euclidean distance is the most common distance measure, though other metrics such as the Manhattan distance can also be used. In order to reduce the affect of units upon the distance measure, it is also common to normalize the features by one of a variety of methods. However, the units seemed to have little effect on the results obtained for the comparison in Section 5. Thus, the Euclidean distance d between a training sample $\boldsymbol{t_m}$ and the test pattern $\boldsymbol{x}$, calculated using

$$d = (\boldsymbol{x} - \boldsymbol{t_m})(\boldsymbol{x} - \boldsymbol{t_m})^T\,, \tag{5}$$

is the distance measure used in this paper. Results for K-nearest neighbor, which classifies the terrain based on the class of the K-nearest training samples, are used for comparison purposes in Section 5. When ties among the K-nearest neighbors occur, the test pattern is assigned to the class with the smallest numerical label.

2.4 Neural Networks

In order to use a neural network for classification purposes, an activation function for each node must be defined as well as a target function or set of target values. Commonly used targets include discrimant functions and a "0/1" target value, which says if the i^{th} target value is 1 and all other targets are 0 then $\boldsymbol{x}$ belongs to class ω_i. The difficulty in training neural networks, however, lies in determining appropriate activation functions for a desired target function and then learning the weights associated with each target function. These weights are typically determined through back propagation which can become

a tedious and difficult process for larger networks. For this reason, as well as slow execution time and reduced popularity in the field of pattern recognition, a neural network classifier is not used in the vibration-based terrain classification comparison.

2.5 Decision Trees

Decision trees attempt to determine a series of rules that determine the expected class of a test pattern. Although decision trees are easy and extremely fast to implement, it can be extremely difficult to determine an appropriate set of rules. This becomes especially true for training sets containing many training samples with a large number of continuous feature variables. As this is the case for vibration-based terrain classification, a decision tree classifier is not considered in the subsequent comparison.

3. FEATURE DESCRIPTION AND CROSS VALIDATION

As is becoming common in vibration-based terrain classification this paper considers frequency domain features from the vertical acceleration $\ddot{z}$, the roll rate ω_{roll} and the pitch rate ω_{pitch}, of the robot. However, unlike previous approaches to terrain classification the vehicle speed v, will also be used as a feature for classification. This eliminates the need for training separate classifiers at each of the considered vehicle speeds. In general, using a speed feature instead of using separate classifiers for each speed creates a slight reduction in accuracy (typically 1%-3%) but for the purposes of this paper, it helps to illustrate the strengths and weaknesses of each classifier. Thus a training pattern $\boldsymbol{t}$ is of the form:

$$\boldsymbol{t} = \left[\boldsymbol{t}_{\ddot{z}}^{T} \quad \boldsymbol{t}_{\omega_{roll}}^{T} \quad \boldsymbol{t}_{\omega_{pitch}}^{T} \quad v\right]^{T}, \tag{6}$$

where v is the vehicle speed and $\boldsymbol{t}_{\omega_{pitch}}^{T}$ is the frequency response magnitude vector of the pitch rate signal ω_{pitch}. Similarly, $\boldsymbol{t}_{\omega_{roll}}^{T}$ and $\boldsymbol{t}_{\ddot{z}}^{T}$ are respectively the frequency response magnitude vectors of the ω_{roll} and $\ddot{z}$ signals. These frequency response magnitudes are computed using a FFT. Consistent with modern vibration-based techniques, PCA is then applied for the purposes of feature extraction, feature selection, and dimensionality reduction. The benefit of PCA is a reduced classification time and the ability to select features that are more appropriate for the chosen classifier. Details on the implementation of PCA as well as its classification uses can be found in (Dupont, et al., 2008a). Since PCA determines the number of features considered for classification, PCA can influence the number of computations used by a given classifier, which in turn affects the classification time. This means that longer classification times can be the result of the necessary PCA energy percentage and not the individual classifiers. Thus the classification times reported in this paper are based on a set PCA energy percentage and not on the energy percentage used to yield the reported mean and minimum accuracies. The classification process uses the test features that have been transformed into eigenspace, which are the result of PCA implementation, and sends them to the classifiers. The classifier then determines the terrain.

In this paper a maximum of two tuning parameters occur for each of the classification schemes. One is the PCA energy percentage and the second, when it exists, is a classifier tuning parameter. The classifier parameters are the width of the window σ for Parzen window estimation, K for K-nearest neighbor, and the degree of the polynomial kernel n or width of the radial basis kernel γ for support vector machines. With only two tuning parameters a brute force search can be reasonably implemented for tuning the algorithm parameters. This tuning process required obtaining the individual class accuracies of the classifier on the cross validation set. Ultimately, parameters that resulted in a maximum product of individual class accuracies are chosen as the best combination as suggested in (Coyle and Collins, 2009).

4. EXPERIMENTAL SET-UP AND DATA COLLECTION

Classification results are based on data recorded from the inertial measurement unit (IMU) on an ATRV-Jr mobile robot. This robot and the IMU are shown in Fig. 1.

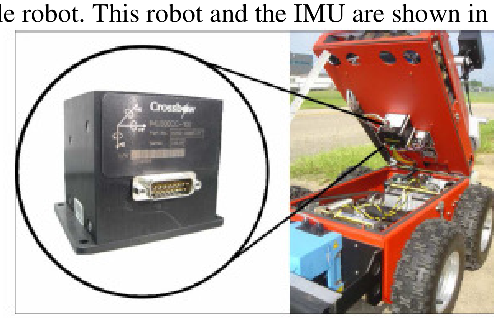

Fig. 1: ATRV-Jr mobile robot and equipped inertial measurement unit

This IMU has the ability to measure the desired signals, vertical acceleration $\ddot{z}$, the roll rate ω_{roll}, and the pitch rate ω_{pitch}. Data was collected by commanding the robot to drive in a straight line for approximately 30 seconds while the desired signals were recorded at 200 Hz. Eight speeds were considered and they are fairly evenly distributed on the interval [0.2 1.4] m/s. These speeds are 0.2, 0.4, 0.5, 0.6, 0.8, 1.0, 1.2, and 1.4 m/s. Seven distinctly different terrains were considered in this data collection: beach sand, packed clay, regular grass, tall grass, loose gravel, packed gravel and asphalt. However, due to the inability to charge the robot's batteries at some

locations, data for some terrains could not be collected at each desired speed. For this reason data was only collected at 0.5 m/s and 1 m/s on beach sand, tall grass, and packed gravel. These two speeds are typical low and high speeds of operation for the ATRV-Jr. As a result, the amount of available samples is not the same for each terrain.

After data collection, half of the trials from each terrain and speed were separated to be used for training. Of the remaining half, one fourth was selected at random to be included in the cross validation set and the rest for testing. This is consistent with suggested levels of cross validation data (Duda et al., 2001). The data was then segmented into one second intervals of time. Each one second time signal was then processed via FFT and PCA as stated in Section 3.

5. EXPERIMENTAL RESULTS

Terrain classification for the purposes of improved vehicle control requires consideration of several important factors. First, high *general* accuracy is needed. This means that not only is high overall accuracy needed, but the algorithm needs to be able to distinguish every possible terrain. Low accuracy of even a single terrain can become problematic. Second, terrain classification must be implemented in real time. This means that algorithms resulting in classification times that are deemed too slow will be impossible to implement with a terrain dependent control system; in general the faster the classification time, the better the algorithm for running online. Lastly, difficulty in training must also be a consideration. Keeping these factors in mind, Table 1 summarizes the performance of the classifiers discussed in Sections 2 and 3 on the ATRV-Jr vibration data collected as described in Section 4.

A SVM with a radial basis kernel function was found to be the most accurate vibration-based classification method in (Weiss et al., 2007), but this experiment showed that a radial basis kernel as well as a polynomial kernel can be highly effective in terms of classification accuracy. This effectiveness is characterized by the highest mean accuracies and second highest minimum accuracies. Also, the test time for SVMs is adequate as it falls somewhere in between the fastest and slowest classifiers. However, SVMs can also require long offline training times, a fact which is consistent with (Weiss et al., 2007). These training times can be on the order of hours or even days depending on processor speeds and the size of the training set $\boldsymbol{T}$. Although both kernel functions performed well in this comparison, it is unlikely that the same kernel function will be the best choice in each case of vibration-based terrain classification problem. This may further complicate training as it may be necessary to consider several choices of kernel functions.

Table 1: Classifier accuracy and classification time

Classifier	Mean Accuracy	Minimum Accuracy	Test Time (msec)
Parzen Window Estimation	81.3%	69.7%	116.0
K-Nearest Neighbor (K=1)	77.9%	50.0%	69.8
Maximum Likelihood Estimation	78.6%	41.3%	0.29
Kesler's Construction	77.3%	55.1%	0.04
SVM radial kernel	83.9%	67.9%	4.1
SVM 25th degree polynomial kernel	83.5%	67.9%	5.1

The best classification time belongs to Kesler's construction. This is the direct result of requiring the fewest online computations. The problem with Kesler's construction appears to be that a separating hyperplane is not always a good choice for a decision boundary between classes. This is why Table 1 shows that Kesler's construction performed more poorly in terms of mean and minimum accuracy. However, if classification time becomes extremely critical, it may be appropriate to use Kesler's construction for vibration-based terrain classification.

Parzen window estimation and K-nearest neighbor have longer classification times than the other classifiers considered. This is the direct result of these classifiers' need to use each training sample for an online calculation, while other classifiers perform offline calculations using the training samples and then save a few important variables that will be used for online computations. This means that if $\boldsymbol{T}$ becomes large, both Parzen window estimation and K-nearest neighbor can become too slow for online implementation. Conversely, despite its problems with classification time, Parzen window estimation is shown to perform well in terms of accuracy, yielding the highest minimum accuracy and second highest mean accuracy. Additionally, Fig. 2 shows less variability in terms of individual class accuracies for Parzen window estimation than for the other classifiers. Thus if the training set is small in nature, Parzen window estimation's classification time deficiency may not be significant enough to outweigh its performance in terms of accuracy.

By assuming the form of the probability distribution, maximum likelihood estimation was able to achieve the

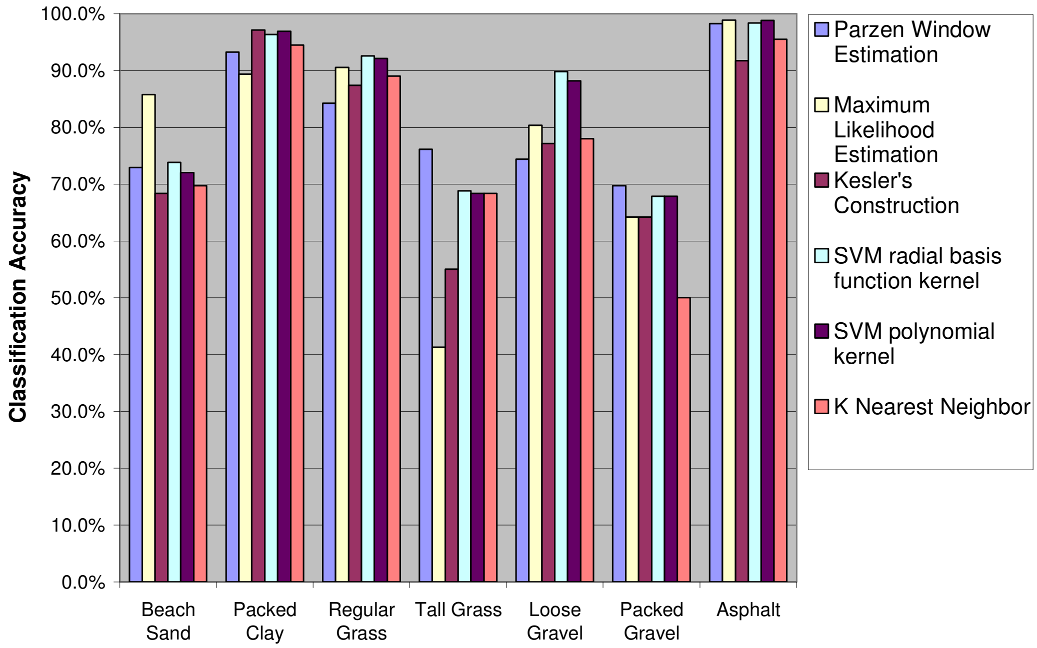

Fig. 2: Classifier performance on individual terrains

second best classification time as well as acceptable mean accuracy. Unfortunately, there is no guarantee that the assumed Gaussian form will even approximately represent the population. This is most likely the reason for its low minimum accuracy of 41.3%, which corresponds to tall grass as seen in Fig. 2. Thus, although maximum likelihood estimation is extremely fast and fairly accurate for many terrains, it is unlikely to perform well on all terrains since Gaussian distributions may not describe the data associated with each terrain.

As expected, these experimental results show that each classifier has both terrain classification strengths and weaknesses. This leads to the possibility of using some sort of hybrid classification technique in order to draw upon the strengths of several types of classifiers. In fact, some new research on staged classification techniques, discussed in (Coyle and Collins, 2009), can be in many ways viewed as a hybrid classifier. This work draws upon the classification time advantages of decision trees and maximum likelihood estimation while using Parzen window estimation and nearest neighbor classifiers to decrease classification error.

6. CONCLUSION

In this paper the strengths and weaknesses of a variety of classification methods have been discussed and demonstrated. Parzen window estimation can be highly affective in terms of accuracy, but the high accuracy comes at the cost of larger classification times. Maximum likelihood estimation is extremely fast to implement but the form of the distribution must be known or closely assumed. Classification time can be a problem for K-nearest neighbor, but its simple implementation and fast training can be of benefit in some cases. Kesler's construction and other forms of linear discriminant functions are extremely fast to implement, but in some cases determining a separating hyperplane can be a problematic endeavor. Support vector machines when paired with the proper kernel function can be affective in achieving high accuracy and adequate classification times. The drawback to support vector machines can be the trial and error process of finding an appropriate kernel function and the extremely long training time that can be associated with tuning the kernel and solving the optimization problem.

As stated in the *No Free Lunch Theorem* (Duda et al., 2001), if the goal is to obtain high accuracy performance, there are no context-independent or usage-independent reasons to favor one learning or classification method over another. This is why this paper has sought to display the context based reasons for favoring one classifier over another for vibration-based terrain classification. It is believed that by looking at the advantages and deficiencies of many statistical classification methods, an appropriate hybrid method can be determined, which can be used for real-time terrain-dependent control systems in the near future.

ACKNOWLEDGEMENT

This work was Prepared through collaborative participation in the Robotics Consortium sponsored by the

U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD 19-01-2- 0012. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. Funding for this research also provided by the National Science Foundation, Project EEC-0540865 and the National Institute of Health 1 R03 HD048465-01.

REFERENCES

Brooks, C., K. Iagnemma, and S. Dubowsky, 2005: Vibration-based Terrain Analysis for Mobile Robots. *IEEE Transactions on Robotics*, **21**, 6 1185-1191.

Chang, C. and C. Lin, 2001: LIBSVM: a library for support vector machines, 2001. [Available online at http://www.csie.ntu.edu.tw/~cjlin/libsvm.]

Coyle, E. and E. G. Collins, Jr. 2009: Vibration-Based Terrain Classification Using Stages and Semi-Automated Tuning. *IEEE International Conf. on Robotics and Automation*, Kobe, Japan, submitted for publication.

Duda, R. O., P. E. Hart, and D. G. Stork, 2001. *Pattern Classification, 2nd Edition,* John Wiley & Sons, 85-90, 164-172, 182-186, 233-234, 259-267, 454, 483-481.

DuPont, E. M., E. G. Collins, E. J. Coyle, R. G. Roberts, 2008: Terrain classification using vibration sensors: theory and methods. *New Research on Mobile Robotics*, Ernest V. Gaines and Lawrence W. Peskov, NOVA, to appear.

DuPont, E. M., C. A. Moore, E. G. Collins, Jr, and E. Coyle. May 2008, Frequency response method for online terrain identification in unmanned ground vehicles. *Autonomous Robots* **24**, 4 337-347.

DuPont, E. M., C. A. Moore, R. G. Roberts, and E. G. Collins, Jr, 2005: Terrain Classification using Probabilistic Neural Networks. *Proc. Of the Florida Conf. on Recent Advances of Robotics*.

DuPont, E. M., C.A. Moore, M. Selekwa, E. Collins, and R. Roberts, November 2005: Online Terrain Classification for Mobile Robots. *Proc. of the ASME International Mechanical Engineering Conference & Exposition*, Orlando, Florida.

DuPont, E. M., R. G. Roberts and C. A. Moore, 2006: The identification of terrains for mobile robots using eigenspace and neural network methods. *Proc. of the Florida Conf. on Recent Advances in Robotics*, Miami, Florida.

Halatci, I., C. Brooks, K. Iagnemma, 2007: Terrain classification and classifier fusion for planetary exploration rovers. *IEEE Aerospace Conf.*

Howard, A., and H. Seraji. 2001: Vision-based terrain characterization and traversability assessment. *Journal of Robotic Systems* **18**, 10 557-587.

Iagnemma, K. D., and S. Dubowsky, 2002: Terrain Estimation for High Speed Rough terrain Autonomous Vehicle Navigation. *Proc. of SPIE Conf. for Unmanned Ground Vehicle Technology IV*, Orlando, FL.

Manduchi, R., A. Castano, A. Talukder, and L. Matthies. 2005: Obstacle detection and terrain classification for autonomous off-road navigation. *Autonomous Robots* **18**, 81-102.

Ojeda, L., J. Borenstein, G. Witus, and R. Karlson, 2006: Terrain Characterization and Classification with a Mobile Robot. *Journal of Field Robotics,* **23**, 3, 103-122.

Sadhukhan, D., and C. Moore. 2003: Online terrain estimation using internal sensors. *Proc. Florida conf. on Recent Advances in Robotics*, Boca Raton, FL.

Se, S., T. Barfoot, P. Jasiobedzki, 2005: Visual motion estimation and terrain modeling for planetary rovers. *Proc. of International Symp. On Artificial Intelligence, Robotics and Automation in Space*, Munich, Germany.

Vandapel, N., M. Hebert, D. Huber and A. Akapuria, 2006: Natural terrain classification using 3-D ladar data for ground robot mobility. *Journal of Field Robotics 23*(10), 839-861.

Vanderwerp, D. February 2005:. What Does Terrain Response Do? [Available online at http://www.caranddriver.com/features/9026/ what-does-terrain-response-do.html].

Weiss, C., N. Fechner, M. Stark, A. Zell, 2007: Comparison of different approaches to vibration-based terrain classification. *European Conf. on Mobile Robotics*, Freiburg, Germany.

Weiss, C., H. Fröhlich, A. Zell, 2006: Vibration-based terrain classification using support vector machines. *International Conf. on Intelligent Robots and Systems*, Beijing, China.

Wolf, D. F., G. S. Sukhatme, D. Fox, and W. Burgard, 2005: Autonomous Terrain Mapping and Classification Using Hidden Markov Models. *Proc. of IEEE International Conf. on Robotics and Automation,* Barcelona, Spain.